arXiv:quant-ph/0112152v1 23 Dec 2001

From Monte Carlo to Quantum Computation

Stefan Heinrich

Fachbereich Informatik
Universität Kaiserslautern
D-67653 Kaiserslautern, Germany
e-mail: heinrich@informatik.uni-kl.de
homepage:
http://www.uni-kl.de/AG-Heinrich

Abstract

Quantum computing was so far mainly concerned with discrete problems. Recently, E. Novak and the author studied quantum algorithms for high dimensional integration and dealt with the question, which advantages quantum computing can bring over classical deterministic or randomized methods for this type of problem.

In this paper we give a short introduction to the basic ideas of quantum computing and survey recent results on high dimensional integration. We discuss connections to the Monte Carlo methology and compare the optimal error rates of quantum algorithms to those of classical deterministic and randomized algorithms.

Introduction

One of the most challenging questions of today, in the overlap of computer science, mathematics, and physics, is the exploration of potential capabilities of quantum computers. Milestones were the algorithm of Shor [26], who showed that quantum computers could factor large integers efficiently (which is widely believed to be infeasible on classical computers) and the quantum search algorithm of Grover [10], which provides a quadratic speedup over deterministic and randomized classical algorithms of searching a database.

So far research was mainly concentrated on discrete problems like the above and many others one encounters in computer science. Much less is known about computational problems of analysis, including such typical field of application of Monte Carlo methods as high dimensional integration. We seek to understand how well these problems can be solved in the quantum model of computation (that is, on a – hypothetical – quantum computer) and how the outcome compares to the efficiency of deterministic or Monte Carlo algorithms on a classical (i. e. non-quantum) computer.

Preprint submitted to Elsevier Science

Abrams and Williams [2] suggested first ideas about quantum integration algorithms. A systematic study was begun by Novak [23], who considered integration of functions from Hölder spaces. This line of research was continued by the author [14], where quantum algorithms for the integration of L_p-functions and, as a key prerequisite, for the computation of the mean of p-summable sequences were constructed. In [14] a rigorous model of quantum computation for numerical problems was developed, as well. The case of integration of functions from Sobolev spaces is considered in [15], and more on the computation of the mean is presented in [17]. These papers also established matching lower bounds. A short survey of first results can be found in [16]. Path integration is studied by Traub and Woźniakowski in [30].

Combining these results with previous ones of information-based complexity theory about the best possible ways of solving the respective problems deterministically or by Monte Carlo on classical computers, we are now in a position to fairly well answer the question where quantum computation can provide a speedup in high dimensional integration and where not. There are cases where quantum algorithms yield an exponential speedup over deterministic algorithms and a quadratic speedup over randomized ones (on classical computers).

Moreover, there is a close connection of quantum algorithms with Monte Carlo: While computations are carried out on superpositions of classical states of qubit systems, and thus in high parallelism, the result can only be accessed through a measurement, which destroys the superposition and outputs any one of the superposed states – with a certain probability. Thus, these algorithms are probabilistic, Monte Carlo, while, on the other hand, completely different laws govern the computation. Nevertheless various Monte Carlo techniques can be put into use to construct quantum algorithms (which then, combined with special quantum techniques, outperform their classical counterparts).

We start with a brief introduction to the ideas of quantum computing. Then we consider the question, what quantum computers could do in numerical (Monte Carlo related) problems and survey recent results on summation and high dimensional integration. We discuss how to use Monte Carlo methodology for the development of quantum algorithms. Complexity issues and comparisons of the potential of deterministic, randomized and quantum algorithms are considered, as well.

For further reading on quantum computation we recommend the surveys by Aharonov [1], Ekert, Hayden, and Inamori [8], Shor [28], and the monographs by Pittenger [24], Gruska [12], and Nielsen and Chuang [21]. For notions and results in information-based complexity theory see the monographs by Traub, Wasilkowski, and Woźniakowski [31] and Novak [22], and the survey Heinrich [13] of the randomized setting.

1 A Short Introduction to Quantum Computing

1.1 History

The first ideas of using quantum devices for computation were expressed at the beginning of the eighties by Manin [18], see also [19], and Feynman [9]. They observed that simulating quantum mechanics on a classical computer is extremly hard, probably infeasible, since it leads to differential equations whose dimensions are exponential in the number of system components. To overcome this, they suggested the idea to simulate quantum mechanics using quantum devices itself.

In 1985, Deutsch [7] developed the rigorous theoretical basis of quantum computation – the model of a quantum Turing machine, which became the so far most serious and still standing challenge to the Turing-Church Thesis (the latter stating that, roughly, every reasonable physical computing device can be simulated with only polynomial increase of resources on a classical Turing machine).

A breakthrough for quantum computing happened in 1994, when Shor [26] showed that efficient factorization of integers would be possible on a quantum computer, which in turn, would mean the possibility of breaking the foremost public key codes like the RSA cryptosystem. Another fundamental contribution was Grover's [10] discovery of an efficient quantum search algorithm in 1996 (we give some further comments on both Shor's and Grover's algorithms later on).

Since then we witness an explosion of efforts, broad research on quantum algorithms for all kinds of (mostly discrete) problems, on quantum cryptography, and quantum information theory. Physicists are intensively working on how to construct quantum computers, that means, finding quantum mechanical systems that can be manipulated to fulfill the abstractly proposed requirements. Systems with a few qubits are already successfully realized in laboratories.

Important forerunners for the development of quantum algorithms for Monte Carlo related, numerical problems were the work of Boyer, Brassard, Høyer, Mosca, and Tapp [4], [5] on quantum counting and the results of Beals, Buhrman, Cleve, Mosca and de Wolf [3] and Nayak and Wu [20] on lower bounds.

In the classical physical world, in classical computation, a bit is represented by two states of a physical system (e.g. charge or no charge) $|0\rangle$, $|1\rangle$. In the (sub)atomic world, which is governed by the laws of quantum mechanics, we have, along with the classical states $|0\rangle$, $|1\rangle$ (such states could be, e.g., spin up or spin down of an electron), also their superpositions:

$$\alpha_0 |0\rangle + \alpha_1 |1\rangle \quad (\alpha_0, \alpha_1 \in \mathbf{C},\ |\alpha_0|^2 + |a_1|^2 = 1),$$

that is, the linear combination of the classical states $|0\rangle$ and $|1\rangle$. The crucial problem however is that if the system is in state $\alpha_0 |0\rangle + \alpha_1 |1\rangle$, we cannot access, measure α_0 and α_1 directly. Instead, a measurement destroys the superposition, causing the system to return to state $|0\rangle$ with probability $|\alpha_0|^2$ and to state $|1\rangle$ with probability $|\alpha_1|^2$.

The mathematical framework for a quantum bit, that is, a quantum system with two classical basis states, is the two dimensional complex Hilbert space $H_1 := \mathbf{C}^2$. Let e_0, e_1 be its unit vector basis. Following quantum mechanical notation, we write $|0\rangle$ instead of e_0 and $|1\rangle$ instead of e_1. The unit sphere of H_1, i.e., the set of all elements of norm 1, is the set of states of the qubit.

A system of m interacting qubits (m-qubit system) is represented by the tensor product

$$H_m := \underbrace{H_1 \otimes H_1 \otimes \ldots \otimes H_1}_{m},$$

which is the 2^m-dimensional complex Hilbert space. It has the canonical basis

$$e_{i_0} \otimes e_{i_1} \otimes \ldots \otimes e_{i_{m-1}} \quad (i_0, i_1, \ldots, i_{m-1}) \in \{0,1\}^m.$$

We make further notational conventions:

$$e_{i_0} \otimes e_{i_1} \otimes \ldots \otimes e_{i_{m-1}} =: |i_0\rangle |i_1\rangle \ldots |i_{m-1}\rangle =: |i\rangle ,$$

where $i := (i_0 i_1 \ldots i_{m-1})_2 := \sum_{k=0}^{m-1} i_k 2^{m-1-k}$. The vectors $|i\rangle$ $(i = 0, \ldots, 2^m - 1)$ are called classical states, while a general state has the form

$$|\psi\rangle = \sum_{i=0}^{2^m-1} \alpha_i |i\rangle \qquad \left(\sum_{i=0}^{2^m-1} |\alpha_i|^2 = 1 \right).$$

As in the case of a single qubit, measuring an m-qubit system in a superposition state $|\psi\rangle$ results in one of the classical states $|i\rangle$ with probability $|\alpha_i|^2$ $(i = 0, \ldots, 2^m - 1)$. So a state is a linear combination of all possible classical states, the coefficients giving the probability that after measurement the system moves to this state.

One more notational convention: if $|\cdot\rangle$ contains a number inside, or a symbol which is used to denote such a number, we mean the canonical basis vector corresponding to this number, if $|\cdot\rangle$ contains just a general symbol, like $|\psi\rangle$, we mean any vector of H_m (this should be clear from the context).

1.3 Quantum Computation

How to use such systems for computing? To make this clear, let us first consider an example of a classical computation – the addition of two m-bit numbers, which we write as follows:

$$\begin{array}{c} |i_0\rangle \ldots |i_{m-1}\rangle\, |j_0\rangle \ldots |j_{m-1}\rangle\, |0\rangle \;\; \ldots \; |0\rangle \\ \downarrow \\ |i_0\rangle \ldots |i_{m-1}\rangle\, |j_0\rangle \ldots |j_{m-1}\rangle\, |k_0\rangle \ldots |k_m\rangle \end{array}$$

This computation is realized using circuits of classical gates (`and, or, not, xor`) in the usual way: add the last bits, then the second last plus the carry bit etc. Let us emphasize here: *Classically, we add two numbers at a time.*

How to operate m-qubit quantum systems? Which operations are allowed? Schrödinger's equation implies: all evolutions of a quantum system must be represented by unitary transforms of H_m. Here is the starting point:

Quantum computing assumes that we are able to perform a number of elementary unitary transforms (quantum gates) on the system.

What are these operations? Let us consider one standard set $\mathcal{G}_m$ of them. First we describe the one-qubit gates – these are gates that manipulate only one component of the tensor product $H_m = H_1 \otimes H_1 \otimes \ldots \otimes H_1$. The Walsh-Hadamard gate $W : H_1 \to H_1$ is defined by

$$W\,|0\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \qquad W\,|1\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

(the values on the basis vectors define the unitary transform uniquely). Its action on the j-th component of H_m is then given by the unitary operator

$$W_m^{(j)} = Id \otimes \ldots \otimes Id \otimes \underbrace{W}_{j} \otimes Id \otimes \ldots \otimes Id,$$

where Id stands for the identity operator on H_1. For a real parameter $0 \le \theta < 2\pi$ the phase shift $P_\theta : H_1 \to H_1$ is defined as

$$P_\theta\,|0\rangle = |0\rangle \qquad P_\theta\,|1\rangle = e^{\imath\theta}\,|1\rangle\,.$$

We define $P^{(j)}_{\theta,m}$ in the respective way. Next we consider two-qubit gates – they manipulate any chosen two components of $H_1 \otimes H_1 \otimes \ldots \otimes H_1$. The quantum xor gate (also called controlled-not gate) $X : H_1 \otimes H_1 \to H_1 \otimes H_1$ is given by

$$\begin{aligned}
X\,|0\rangle\,|0\rangle &= |0\rangle\,|0\rangle \\
X\,|0\rangle\,|1\rangle &= |0\rangle\,|1\rangle \\
X\,|1\rangle\,|0\rangle &= |1\rangle\,|1\rangle \\
X\,|1\rangle\,|1\rangle &= |1\rangle\,|0\rangle
\end{aligned}$$

That is, if the first bit is zero, nothing happens to the second, and if the first is one, the second is negated (controlled not). We can look at this gate also as follows: the xor of the two bits replaces the second bit. Denote by $X^{(k,\ell)}_m : H_m \to H_m$ the unitary operator given by applying X to the k-th and ℓ-th component, that is,

$$X^{(k,\ell)}_m\,|i_0\rangle \ldots |i_k\rangle \ldots |i_\ell\rangle \ldots |i_{m-1}\rangle \;= |i_0\rangle \ldots |y\rangle \ldots |z\rangle \ldots |i_{m-1}\rangle\,,$$

where

$$|y\rangle\,|z\rangle := X\,|i_k\rangle\,|i_\ell\rangle\,.$$

Now we define

$$\mathcal{G}_m = \Big\{W^{(j)}_m,\, P^{(j)}_{\theta,m},\, X^{(k,\ell)}_m \;:\; 0 \le j, k \ne \ell \le m-1,\, 0 \le \theta < 2\pi\Big\}\,.$$

The following two results can be found, e.g., in [21].

Theorem 1 *The set $\mathcal{G}_m$ is a universal system of gates – each unitary transform of H_m can be represented as a finite composition of elements of $\mathcal{G}_m$ (up to a complex scalar factor).*

The set $\mathcal{G}_m$ is still an infinite set. Now consider the following finite subset:

$$\mathcal{G}^0_m = \Big\{W^{(j)}_m,\, P^{(j)}_{\pi/4,m},\, X^{(k,\ell)}_m \;:\; 0 \le j, k \ne \ell \le m-1\Big\}\,.$$

Theorem 2 *The set $\mathcal{G}^0_m$ forms an approximately universal system of gates – each unitary transform of H_m can be approximated in the operator norm to each precision by a finite composition of elements of $\mathcal{G}^0_m$ (again up to a complex factor).*

So once we can implement these gates we can do all unitary transforms (of course, the efficiency of such an approximation is still an issue, we will come back to that later on). Physicists are working on implementations of these gates in various quantum systems such as photons, trapped ions, magnetic resonance systems etc. Let me emphasize some crucial points:

1. These gates can transform classical states into superpositions. For example, the Hadamard gate applied to the first and then to the second qubit:

$$|0\rangle\,|0\rangle \longrightarrow \frac{1}{2}\left(|0\rangle\,|0\rangle + |0\rangle\,|1\rangle + |1\rangle\,|0\rangle + |1\rangle\,|1\rangle\right).$$

2. They also act on superpositions. For instance:
2.1 The quantum xor:

$$\begin{array}{c} \alpha_0\,|0\rangle\,|0\rangle + \alpha_1\,|0\rangle\,|1\rangle + \alpha_2\,|1\rangle\,|0\rangle + \alpha_3\,|1\rangle\,|1\rangle \\ \downarrow \\ \alpha_0\,|0\rangle\,|0\rangle + \alpha_1\,|0\rangle\,|1\rangle + \alpha_2\,|1\rangle\,|1\rangle + \alpha_3\,|1\rangle\,|0\rangle \end{array}$$

2.2 Quantum addition of binary numbers:

$$\begin{array}{c} \sum \alpha_{ij}\,|i_0\rangle \ldots |i_{m-1}\rangle\,|j_0\rangle \ldots |j_{m-1}\rangle\,|0\rangle \;\; \ldots \;|0\rangle \\ \downarrow \\ \sum \alpha_{ij}\,|i_0\rangle \ldots |i_{m-1}\rangle\,|j_0\rangle \ldots |j_{m-1}\rangle\,|k_0\rangle \ldots |k_m\rangle \end{array}$$

That is, in the quantum world, we add all possible binary m-digit numbers in parallel (assuming that all $\alpha_{ij} \neq 0$). So, is a quantum computer an ideal parallel computer, with exponentially many processors? Not exactly, it is not that easy! We cannot access all components of the superposition. We have to measure, thus destroying all results but one! We would get

$$|i_0\rangle \ldots |i_{m-1}\rangle\,|j_0\rangle \ldots |j_{m-1}\rangle\,|k_0\rangle \ldots |k_m\rangle$$

with probability $|\alpha_{ij}|^2$. In this sense a quantum computer is like a black box – we can manipulate it, but we cannot "look into it". Anyway, we have reached the quantum model of computation (the general way a quantum algorithm should look like):

Quantum model of computation:

starting state: $|i_0\rangle \in H_m$, a classical state

computation: $U_1, \ldots, U_n \in \mathcal{G}_m$

$|i_0\rangle \to U_1\,|i_0\rangle \to U_2U_1\,|i_0\rangle \ldots$

$\to |\psi\rangle := U_nU_{n-1}\ldots U_2U_1\,|i_0\rangle$

measurement: $|\psi\rangle = \sum_{i=0}^{2^m-1} \alpha_i\,|i\rangle$

$\to |i\rangle$ with probability $|\alpha_i|^2$

repeat

The cost of the algorithm is defined to be the number of gates n. We may allow for a classical computation after each measurement to determine the next starting state, and, after the last measurement, the final result (see [14] for a formal approach). Then, speaking about the cost of the algorithm, we also have to take into account the needed classical gates. A quantum algorithm is said to solve a given problem if it produces a solution with probability $\geq 3/4$ (any fixed constant $\gamma > 1/2$ would do, since the success probability can always be increased by suitable repetitions).

1.4 Computational Power of Quantum Computers

What can we do with a quantum computer that we cannot do on a classical computer? A simple answer related to Monte Carlo algorithms can be given as follows: On classical computers, pseudo-random numbers (that is, in fact, fully deterministic numbers) are used instead of random numbers. A quantum computer would provide us with true randomness. It is simple to let a quantum computer act as a random number generator:

Quantum random number generator

$$
\begin{array}{ll}
|0\rangle \ldots |0\rangle & \\
\quad\downarrow & \text{(Walsh-Hadamard transform)} \\
2^{-m/2} \sum_{i=0}^{2^m-1} |i\rangle & \\
\quad\downarrow & \text{(measurement)} \\
|i\rangle \text{ with probability } 2^{-m} &
\end{array}
$$

So, on a quantum computer, we can implement Monte Carlo with true (physically based) randomness. Let us go one step further: What can be done on a quantum computer that a *randomized* classical computer (that is, a classical computer with access to true randomness) cannot do? We already mentioned Shor's result [26]: An algorithm which, for a composite integer N, finds a nontrivial factor with cost $\mathcal{O}((\log N)^3)$. No polynomial in $\log N$ classical (deterministic or randomized) algorithm is known. Recall that there are polynomial randomized prime number tests (see Solovay and Strassen [29], and Rabin [25]), but they don't give the factors when they find compositeness.

Although not very likely, it is not excluded that there are polynomial classical factoring algorithms (maybe even with cost $\mathcal{O}((\log N)^3)$). So formally, the above result does not strictly prove the superiority of quantum computing.

Matters are different with Grover's result [10]. The problem is the following: Given $f : \{0, \ldots, N-1\} \to \{0, 1\}$ with the property that there is a unique i_0 with $f(i_0) = 1$, find this i_0. It is not difficult to see that classical (deterministic or randomized) algorithms need cost $\Omega(N)$. Grover [10] produced a quantum algorithm that solves the problem with cost $\mathcal{O}(\sqrt{N})$. Here we have to extend the above model of computation by what is called a quantum query (quantum subroutine, quantum oracle), which is given by the unitary map $Q_f : H_m \to H_m$ defined as

$$Q_f : |i\rangle\,|y\rangle \to |i\rangle\,|y \oplus f(i)\rangle \quad (0 \le i < 2^{m-1},\ y \in \{0, 1\}),$$

where $\oplus$ is addition modulo 2, and we assumed $N = 2^{m-1}$. Note that if the last qubit is set to zero, Q_f is like a subroutine which writes the function value to the last entry. This is what we essentially assume in the design of most numerical and Monte Carlo algorithms: the function we want to handle, e.g., integrate, is given as a subroutine. (In the quantum setting, addition modulo 2 guarantees that Q_f is a bijection on the classical states and thus, extends to a unitary operator.) It can be shown that such a quantum query can always be implemented efficiently (with only a constant factor increase of resources), once we have a classical implementation for f. A quantum computation is then still given by the scheme in section 1.3, but with the difference that $U_i \in \mathcal{S}_m \cup \{Q_f\}$. It is important to mention that the rest of the U_i should not depend on f, that is, the algorithm gets its information about f only through query calls. The cost is defined to be the number of query calls, multiplied by the number of qubits, plus the number of gates .

2 Summation and Integration

Now we consider numerical problems. Let us formulate the general integration problem: Given a set D, a measure μ on it, and a μ-integrable function $f : D \to \mathbf{R}$, compute (approximate)

$$\int_D f(t)d\mu(t).$$

As its simplest instant, this includes computing the mean (and hence the sum) of a finite sequence:

$$D = \{0, \ldots, N-1\}, \quad f : D \to \mathbf{R}, \quad \mu(\{i\}) = \frac{1}{N}$$

$$\int_D f(t)d\mu(t) = \frac{1}{N} \sum_{i=0}^{N-1} f(i).$$

Theorem 3 (Brassard, Høyer, Mosca, and Tapp [5]) *For all n, N, $n < N$, there is a quantum algorithm which computes the mean $\frac{1}{N}\sum_{i=0}^{N-1} f(i)$ for all sequences $f(i)$ with $f(i) \in \{0,1\}$ with cost $\widetilde{\mathcal{O}}(n)$ and error $\mathcal{O}(1/n)$.*

We use the $\widetilde{\mathcal{O}}$ notation to indicate that we suppress possible logarithmic factors. The result is easily extended to real valued $f(i) \in [-1,1]$. We have to specify how a quantum algorithm can access such values of f. Without loss of generality we suppose that N is of the form $N = 2^{m_1}$. We assume that we have a quantum query (subroutine) providing essentially the first m_2 digits of $f(i)$ for some suitable m_2. More precisely, we can use the unitary mapping Q_f on $H_m = H_{m_1} \otimes H_{m_2}$ defined as

$$Q_f : |i\rangle\,|y\rangle \to |i\rangle\,|y \oplus [f(i)]_{m_2}\rangle \quad (0 \le i < 2^{m_1},\, 0 \le y < 2^{m_2}).$$

Here $\oplus$ is addition modulo 2^{m_2} and $[f(i)]_{m_2}$ is the m_2-bit integer $\lfloor 2^{m_2-1}(f(i)+1)\rfloor$ if $f(i) < 1$ and $2^{m_2}-1$ if $f(i) = 1$. (If we have to deal with functions taking values in other intervals, say $[a,b]$, we use a similar encoding as a binary integer, $[f(i)]_{m_2} = \lfloor 2^{m_2}(f(i)-a)/(b-a)\rfloor$ if $f(i) < b$ and $[f(i)]_{m_2} = 2^{m_2}-1$ if $f(i) = b$.) We say that a quantum algorithm computes a number with error ε, if with probability $\ge 3/4$ the result of the algorithm is within ε of that number. For more details on queries for numerical problems and the error definition we refer to [14]. Theorem 3 should be viewed under the aspect of huge N and moderate n. For comparison, let us mention that for classical deterministic algorithms the error at cost n for $n < N/2$ is $\Omega(1)$ (even $n < cN$ for any fixed $0 < c < 1$ suffices). In the classical randomized setting we obtain $\Theta(1/\sqrt{n})$ (see [16] for details on the latter two statements).

How to treat integration? Define the Hölder classes $\mathcal{F}_d^{r,\varrho}$ for $r \in \mathbf{N}_0$, $d \in \mathbf{N}$, and $0 < \varrho \le 1$ as

$$\begin{aligned}\mathcal{F}_d^{r,\varrho} = \{ & f \in C^r([0,1]^d) \,:\, \|f\|_\infty \le 1, \\ & |\partial^\alpha f(x) - \partial^\alpha f(y)| \le |x-y|^\varrho, |\alpha| = r\}.\end{aligned}$$

Here $C^r([0,1]^d)$ stands for the set of r times continuously differentiable functions on $[0,1]^d$, $\|\ \|_\infty$ denotes the supremum norm, α represents a multiindex, and ∂^α is the respective partial derivative.

Theorem 4 (Novak [23]) *For all n, there is a quantum algorithm which computes the integral of functions $f \in \mathcal{F}_d^{r,\varrho}$ with cost $\widetilde{\mathcal{O}}(n)$ and error $\mathcal{O}(n^{-(r+\varrho)/d-1})$.*

Because of its relation to Monte Carlo methodology it is interesting to describe the idea behind the algorithm: We use the technique of separation of the main part (in other words, of introducing a control variate) to reduce integration to summation in an optimal way. Let P_n be some interpolation operator with

n nodes $(\tau_i)_{i=0}^{n-1} \in D = [0,1]^d$ suitable for $\mathcal{F}_d^{r,\varrho}$ in the sense that it gives the optimal approximation order

$$\sup_{t\in D} |f(t) - P_n f(t)| = \mathcal{O}(n^{-(r+\varrho)/d}).$$

Represent

$$\int_D f(t)dt = \int_D P_n f(t)dt + \int_D (f(t) - P_n f(t))dt. \tag{1}$$

Clearly,

$$\int_D P_n f(t)dt = \sum_{i=0}^{n-1} c_i f(\tau_i)$$

is a quadrature, which can be computed exactly, classically, with $\mathcal{O}(n)$ cost. The second part is approximated suitably (see [23] for details) in the form

$$\int_D (f(t) - P_n f(t))dt \approx \frac{1}{N} \sum_{i=0}^{N-1} (f(t_i) - P_n f(t_i)),$$

where N may be much larger than n. Finally we use quantum computation to approximate the mean on the right hand side. It can be shown that the order of the error is the product of the following contributions:

$$\underbrace{n^{-(r+\varrho)/d}}_{\text{separation of main part}} \quad * \quad \underbrace{n^{-1}}_{\text{quantum computation}}$$

Note that if the last term of (1) were approximated by standard Monte Carlo, we would get the exponent $-(r+\varrho)/d - 1/2$.

Compare this result with the classical deterministic setting, where the best we can achieve is an error $\mathcal{O}(n^{-(r+\varrho)/d})$ at cost $\mathcal{O}(n)$ and with the classical randomized setting with optimal error $\mathcal{O}(n^{-(r+\varrho)/d-1/2})$. Look at the exponents when $(r+\varrho)/d$ is small! It follows that the speedup over the deterministic setting can be polynomial with arbitrarily large power. Note that in [14], Section 6, a class of functions of low smoothness is considered and it is proved that for this class the gain of quantum over classical deterministic computation is even exponential. A similar situation is observed in [30] for path integration.

A crucial assumption for the summation algorithm in Theorem 3, and thus also for Theorem 4, is the uniform boundedness of the set of sequences (with bounds independent of n and N). This raised the problem whether the quantum gain was tied to that assumption. In particular, what happens in the case of square summability, which is the most important case for Monte Carlo, that is, $f \in \mathcal{B}_2^N$ or $f \in \mathcal{W}_{2,d}^r$ (see the definitions below)? Will quantum computation retain

its superiority or will Monte Carlo catch up? Even more generally, what about $f \in \mathcal{B}_p^N$ or $f \in \mathcal{W}_{p,d}^r$ for $1 \le p < \infty$? Let us first look at the discrete problem. Denote

$$\mathcal{B}_p^N = \left\{ f \,:\, \frac{1}{N} \sum_{i=0}^{N-1} |f(i)|^p \le 1 \right\}.$$

The summation problem for this class was settled in [14] (case (i) and (ii)) and in [17] (case (iii)). The latter answered a question posed in [16].

Theorem 5 *Let $1 \le p < \infty$. For all n, N, $n < N$, there is a quantum algorithm which computes the mean $\frac{1}{N}\sum_{i=0}^{N-1} f(i)$ for all sequences $f \in \mathcal{B}_p^N$ with cost $\widetilde{\mathcal{O}}(n)$ and error*

(i)	$\mathcal{O}(n^{-1})$	*if* $2 \le p < \infty$,
(ii)	$\mathcal{O}(n^{-2+2/p})$	*if* $1 \le p < 2$ *and* $n < \sqrt{N}$,
(iii)	$\mathcal{O}\left(n^{-2/p} N^{2/p-1}\right)$	*if* $1 \le p < 2$ *and* $\sqrt{N} \le n < N$.

Again, a comparison to the two classical settings might be illustrative (see [16] for details): In the classical deterministic case, for $n < N/2$, nothing better than $\Omega(1)$ can be obtained, in the classical randomized case the optimal rates are $\Theta(n^{-1/2})$ if $2 \le p < \infty$, and $\Theta(n^{-1+1/p})$ if $1 \le p < 2$. Now we consider integration in the Sobolev classes $\mathcal{W}_{p,d}^r$, which are defined by

$$\mathcal{W}_{p,d}^r = \{ f \in L_p([0,1]^d) \,:\, \|\partial^\alpha f\|_{L_p} \le 1,\ |\alpha| \le r \},$$

where $r \in \mathbf{N}$, $1 \le p \le \infty$, and ∂^α is the weak partial derivative. The following result is proved in [15]. It answers another question from [16]. The approach consists of a new discretization technique, by which one can derive (optimal) integration algorithms for $\mathcal{W}_{p,d}^r$ from (optimal) summation algorithms for $\mathcal{B}_p^N$.

Theorem 6 *Let $1 \le p < \infty$, $r, d \in \mathbf{N}$, $r/d > 1/p$ (Sobolev embedding condition). For all n, there is a quantum algorithm which computes the integral of functions $f \in \mathcal{W}_{p,d}^r$ with cost $\widetilde{\mathcal{O}}(n)$ and error $\mathcal{O}(n^{-r/d-1})$.*

In the classical deterministic setting the optimal rate is $\mathcal{O}(n^{-r/d})$, while in the classical randomized setting we have $\mathcal{O}(n^{-r/d-1/2})$ if $2 \le p < \infty$, and $\mathcal{O}(n^{-r/d-1+1/p})$ if $1 \le p < 2$ (see [16], [13], and the references therein). The quantum rate for $1 \le p < 2$ comes as a surprise. After previous results one was tempted to conjecture that the quantum setting could reduce the exponent of the classical randomized setting by at most $1/2$. Now we see ($p = 1$) that there can even be a reduction by 1.

Are these results about quantum algorithms optimal? In other words, is it possible to improve the rates by other, better quantum algorithms? To verify

optimality, we have to establish lower bounds valid for all possible quantum algorithms. It turns out that all the results about summation and integration presented here are optimal (up to logarithmic factors, at least). The following was the first matching lower bound for summation. It shows that the rate in Theorem 3 is optimal.

Theorem 7 (Nayak and Wu [20]) *There are constants $c_1, c_2 > 0$ such that for all n, N with $n < c_1 N$ the following holds: Each quantum algorithm which computes the mean $\frac{1}{N}\sum_{i=0}^{N-1} f(i)$ for all sequences $f(i)$ with $|f(i)| \leq 1$ using at most n quantum queries has error not smaller than c_2/n.*

Using their technique and methods of information-based complexity theory, the following can be shown (see [23], [14], [17], [15]).

Theorem 8 *The rates established in Theorem 4 (integration in $\mathcal{F}_d^{r,\varrho}$), Theorem 5 (mean of sequences in $\mathcal{B}_p^N$), and Theorem 6 (integration in $\mathcal{W}_{p,d}^r$) are optimal (up to logarithmic factors).*

In the following table we summarize the results. The respective entries give the optimal rates at cost n, constants and logarithmic factors are suppressed. The constant $0 < c < 1$ in the first column does not depend on n and N.

	deterministic	random	quantum
$\mathcal{B}_p^N$, $2 \leq p \leq \infty$, $n < cN$	1	$n^{-1/2}$	n^{-1}
$\mathcal{B}_p^N$, $1 < p < 2$, $n < \sqrt{N}$	1	$n^{-1+1/p}$	$n^{-2+2/p}$
$\mathcal{B}_p^N$, $1 < p < 2$, $\sqrt{N} \leq n < cN$	1	$n^{-1+1/p}$	$n^{-2/p}N^{2/p-1}$
$\mathcal{B}_1^N$, $n < \sqrt{N}$	1	1	1
$\mathcal{B}_1^N$, $\sqrt{N} \leq n < cN$	1	1	$n^{-2}N$
$\mathcal{F}_d^{r,\varrho}$	$n^{-(r+\varrho)/d}$	$n^{-(r+\varrho)/d-1/2}$	$n^{-(r+\varrho)/d-1}$
$\mathcal{W}_{p,d}^r$, $2 \leq p \leq \infty$	$n^{-r/d}$	$n^{-r/d-1/2}$	$n^{-r/d-1}$
$\mathcal{W}_{p,d}^r$, $1 < p < 2$	$n^{-r/d}$	$n^{-r/d-1+1/p}$	$n^{-r/d-1}$
$\mathcal{W}_{1,d}^r$	$n^{-r/d}$	$n^{-r/d}$	$n^{-r/d-1}$

References

[1] D. Aharonov (1998): Quantum computation – a review. In: Annual Review of Computational Physics, World Scientific, volume VI, ed. Dietrich Stauffer, see also http://arXiv.org/abs/quant-ph/9812037.

[2] D. S. Abrams and C. P. Williams (1999): Fast quantum algorithms for numerical integrals and stochastic processes. Technical report, http://arXiv.org/abs/quant-ph/9908083.

[3] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf (1998): Quantum lower bounds by polynomials. Proceedings of 39th IEEE FOCS, 352-361, see also http://arXiv.org/abs/quant-ph/9802049.

[4] M. Boyer, P. Brassard, P. Høyer, and A. Tapp (1998): Tight bounds on quantum searching. Fortschritte der Physik **46**, 493 – 505, see also http://arXiv.org/abs/quant-ph/9605034.

[5] G. Brassard, P. Høyer, M. Mosca, and A. Tapp (2000): Quantum amplitude amplification and estimation. Technical report, http://arXiv.org/abs/quant-ph/0005055.

[6] G. Brassard, P. Høyer, and A. Tapp (1998): Quantum counting. Lect. Notes in Comp. Science **1443**, 820 – 831, see also http://arXiv.org/abs/quant-ph/9805082.

[7] D. Deutsch (1985): Quantum theory, the Church-Turing principle and the universal quantum computer. Proc. R. Soc. Lond., Ser. A 400, 97-117.

[8] A. Ekert, P. Hayden, and H. Inamori (2000): Basic concepts in quantum computation. See http://arXiv.org/abs/quant-ph/0011013.

[9] R. Feynman (1982): Simulating physics with computers. Int. J. Theor. Phys. **21**, 467–488.

[10] L. Grover (1996): A fast quantum mechanical algorithm for database search. Proc. 28 Annual ACM Symp. on the Theory of Computing, 212–219, ACM Press New York. See also http://arXiv.org/abs/quant-ph/9605043.

[11] L. Grover (1998): A framework for fast quantum mechanical algorithms. Proc. 30 Annual ACM Symp. on the Theory of Computing, 53–62, ACM Press New York. See also http://arXiv.org/abs/quant-ph/9711043.

[12] J. Gruska (1999): Quantum Computing. McGraw-Hill, London.

[13] S. Heinrich (1993): Random approximation in numerical analysis. In: K. D. Bierstedt, A. Pietsch, W. M. Ruess, and D. Vogt, editors, Functional Analysis, 123 – 171, Marcel Dekker.

[14] S. Heinrich (2001): Quantum summation with an application to integration. Journal of Complexity (to appear). See also http://arXiv.org/abs/quant-ph/0105116.

[15] S. Heinrich (2001): Quantum integration in Sobolev classes (in preparation).

[16] S. Heinrich and E. Novak (2001a): Optimal summation and integration by deterministic, randomized, and quantum algorithms. In: K.-T. Fang, F. J. Hickernell, and H. Niederreiter, editors, Monte Carlo and Quasi-Monte Carlo Methods 2000, Springer-Verlag, Berlin (to appear), see also http://arXiv.org/abs/quant-ph/0105114.

[17] S. Heinrich and E. Novak (2001b): On a problem in quantum summation. Submitted to J. Complexity, see also http://arXiv.org/abs/quant-ph/0109038.

[18] Yu. I. Manin (1980): Computable and uncomputable (in Russian). Sovetskoye Radio, Moscow.

[19] Yu. I. Manin (1999): Classical computing, quantum computing, and Shor's factoring algorithm. See http://arXiv.org/abs/quant-ph/9903008.

[20] A. Nayak and F. Wu (1999): The quantum query complexity of approximating the median and related statistics. STOC, May 1999, 384–393, see also http://arXiv.org/abs/quant-ph/9804066.

[21] M. A. Nielsen and I. L. Chuang (2000): Quantum Computation and Quantum Information. Cambridge University Press.

[22] E. Novak (1988): Deterministic and Stochastic Error Bounds in Numerical Analysis. Lecture Notes in Mathematics **1349**, Springer.

[23] E. Novak (2001): Quantum complexity of integration. J. Complexity **17**, 2–16. See also http://arXiv.org/abs/quant-ph/0008124.

[24] A. O. Pittenger (1999): Introduction to Quantum Computing Algorithms. Birkhäuser, Boston.

[25] M. O. Rabin (1980): Probabilistic algorithm for testing primality. J. Number Theory 12, 128-138.

[26] P. W. Shor (1994): Algorithms for quantum computation: Discrete logarithms and factoring. Proceedings of the 35th Annual Symposium on Foundations of Computer Science, IEEE Computer Society Press, Los Alamitos, CA, pp. 124–134. See also http://arXiv.org/abs/quant-ph/9508027.

[27] P. W. Shor (1998): Quantum computing. Documenta Mathematica, Extra Volume ICM 1998, I, 467–486.

[28] P. W. Shor (2000): Introduction to quantum algorithms. See http://arXiv.org/abs/quant-ph/0005003.

[29] R. Solovay and V. Strassen (1977): A fast Monte-Carlo test for primality. SIAM J. Comput. 6, 84-85.

[30] J. F. Traub and H. Woźniakowski (2001): Path integration on a quantum computer. See http://arXiv.org/abs/quant-ph/0109113.

[31] J. F. Traub, G. W. Wasilkowski, and H. Woźniakowski (1988): Information-Based Complexity. Academic Press.